



## *CD FY10 Budget and Tactical Plan Review*

# **FY10 Tactical Plans for** **GRID**

FermiGrid	CD-3330
Grid Services	CD-3335
Open Science Grid @ FNAL	CD-3326

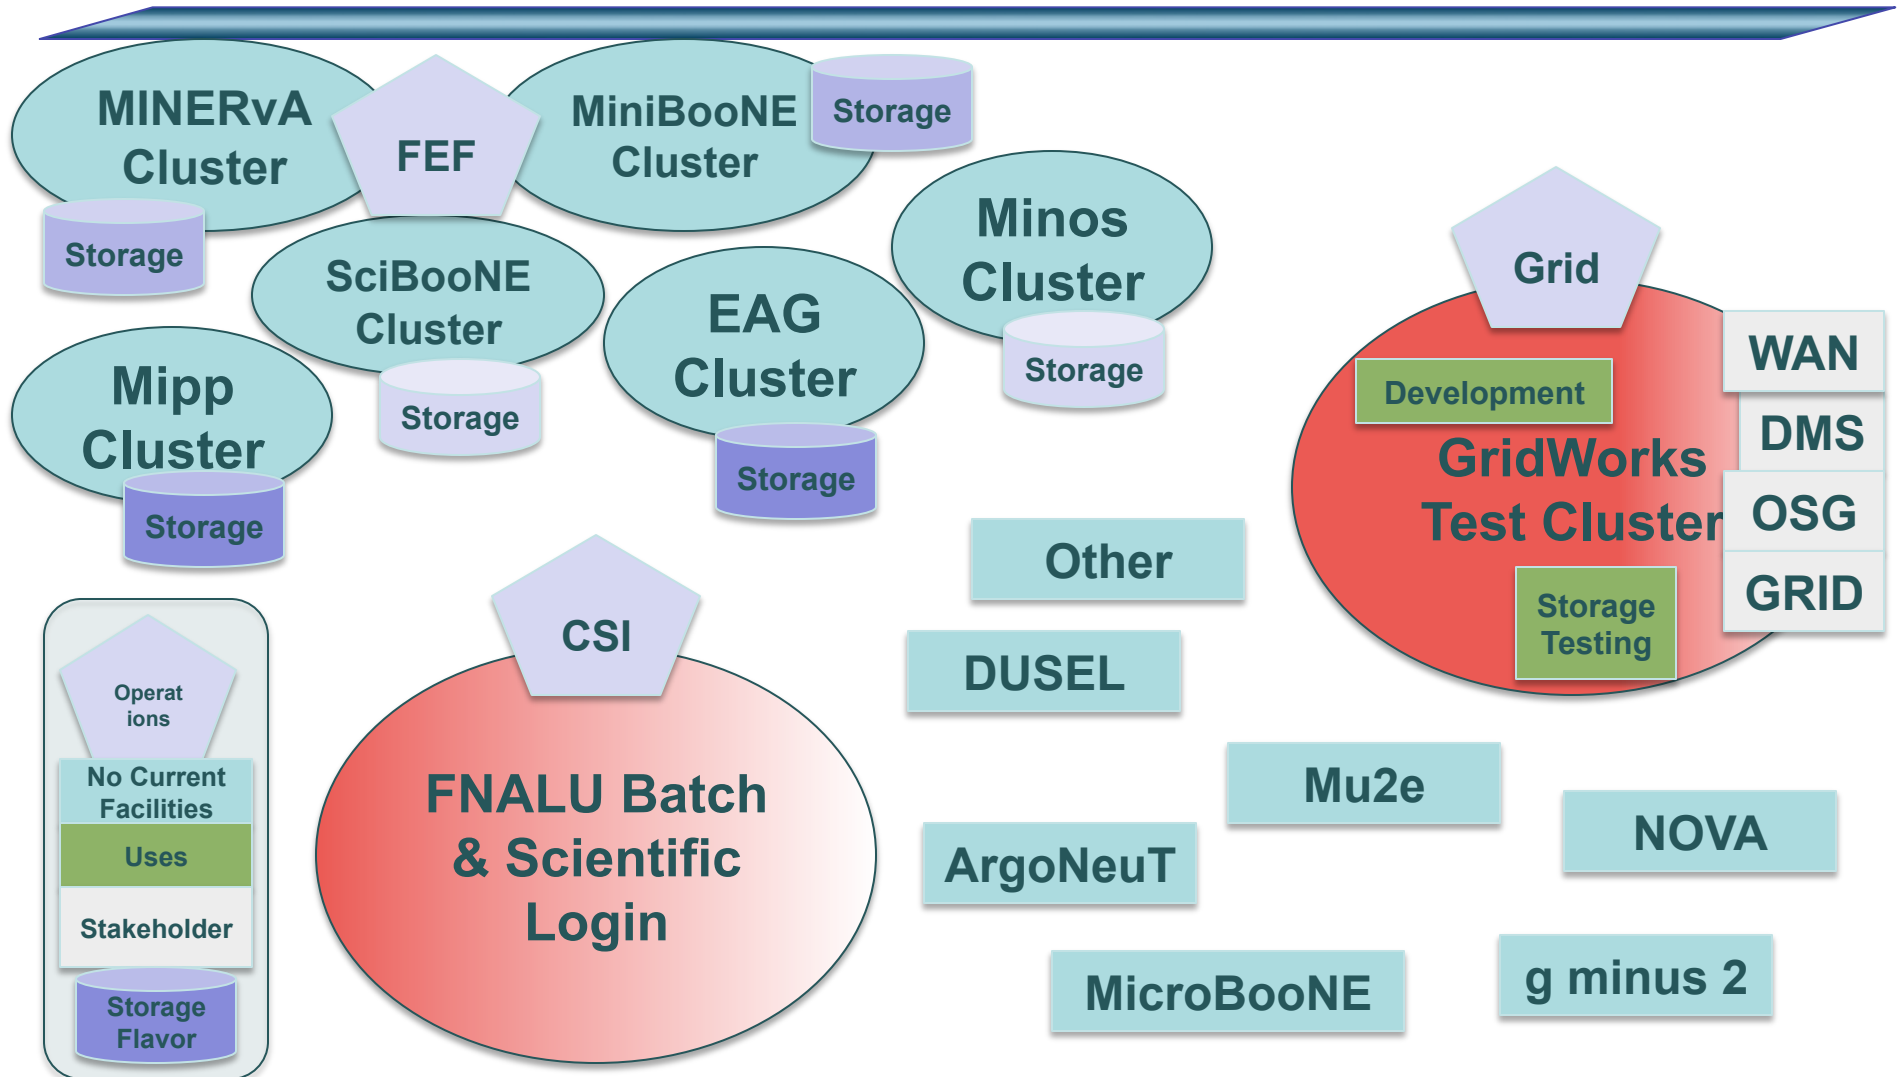
Eileen Berman, Keith Chadwick, Gabriele Garzoglio, Ruth  
Pordes

October 2, 2009

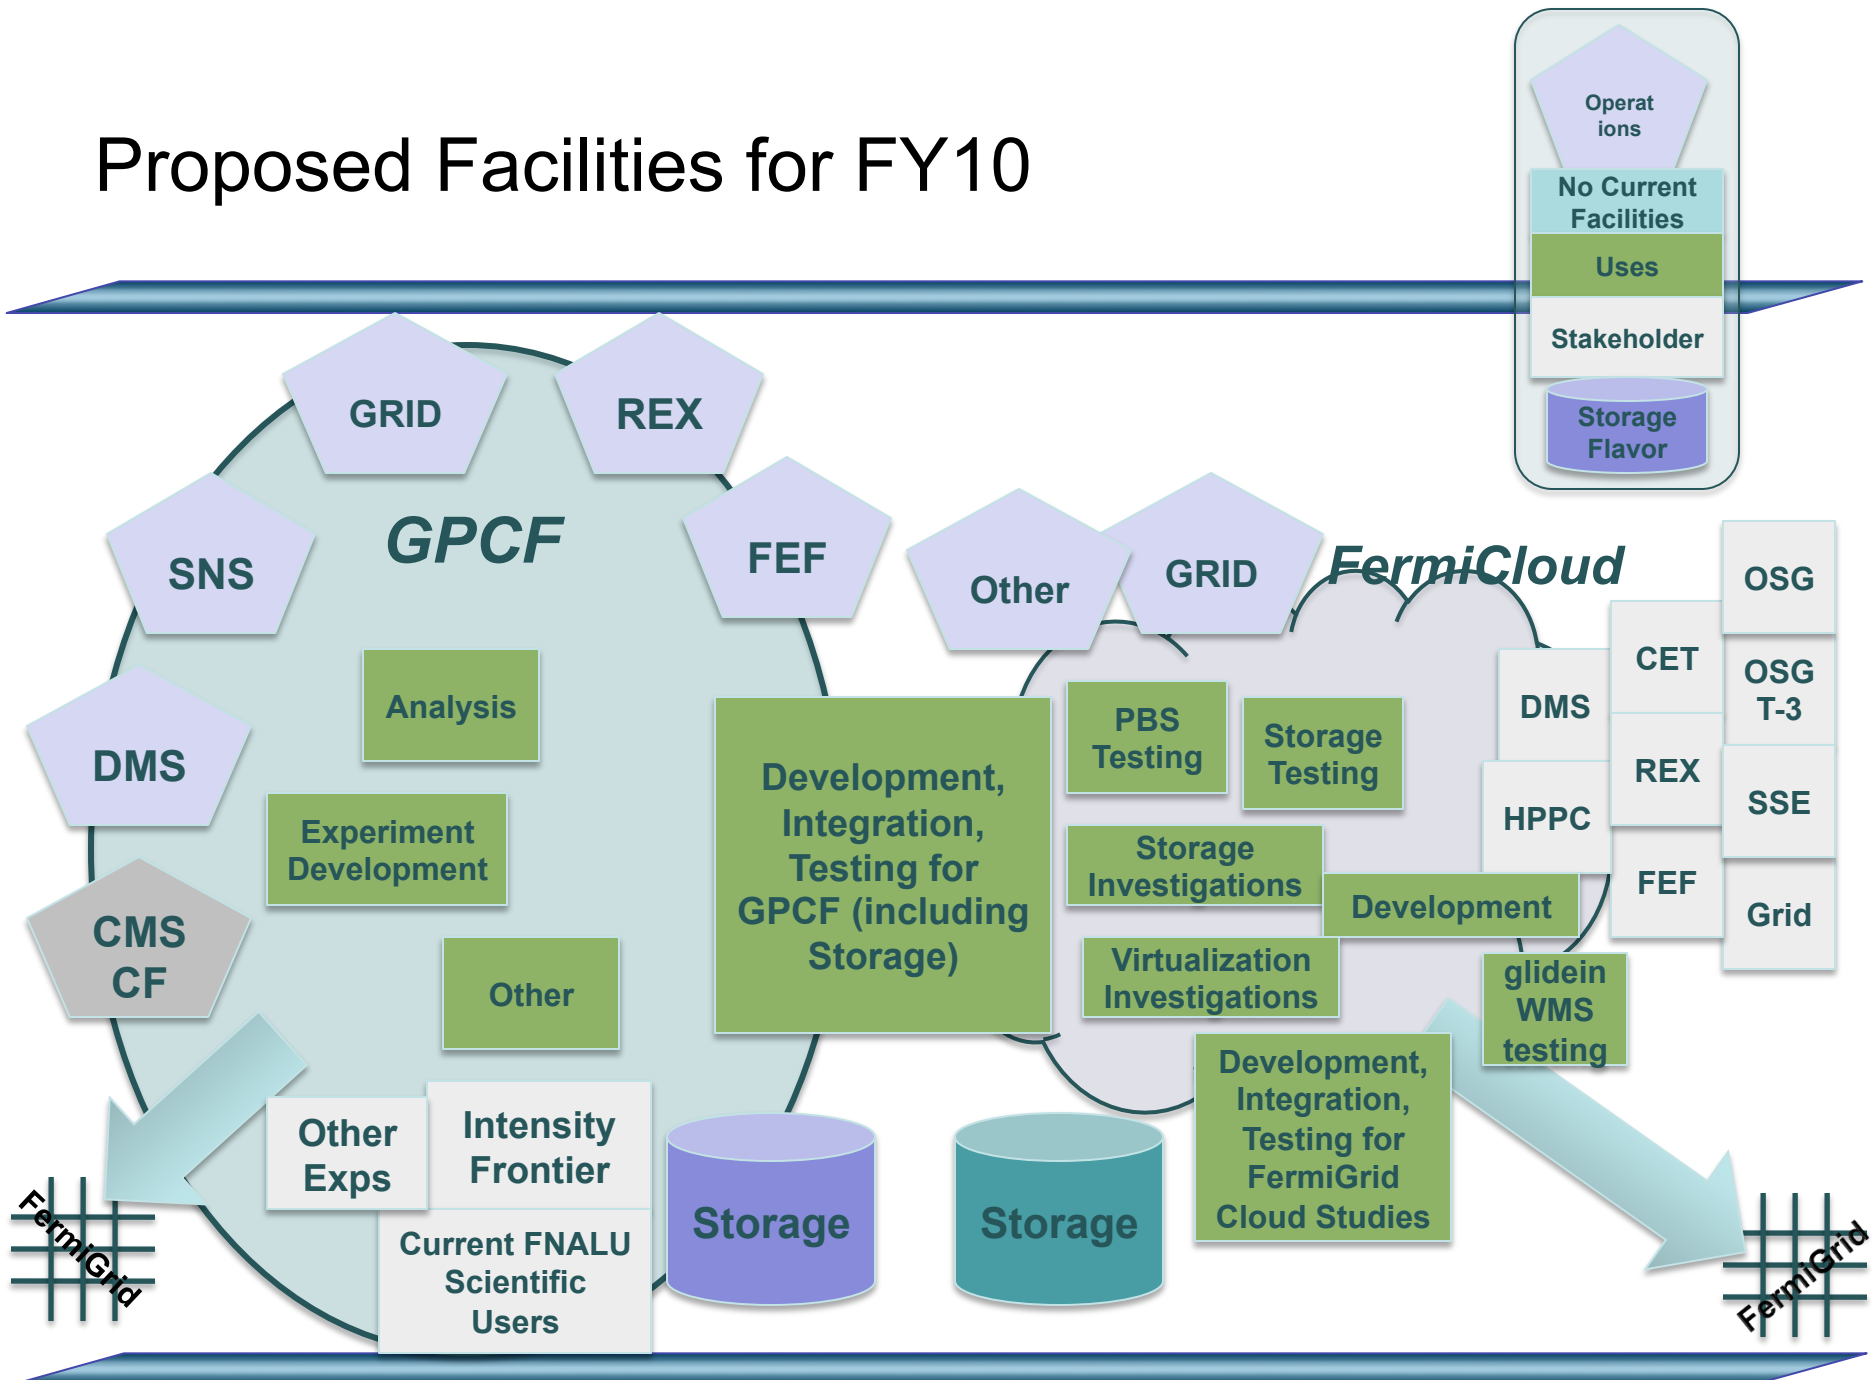
# Introduction – New Grid Activities and Their Context

- FermiCloud – Infrastructure-as-a-Service, on demand provisioning of scalable resources for development, integration, and testing.
- Storage Investigations – Understanding of grid accessible storage use patterns, needs, and solutions. (FermiGrid, FermiCloud, General Physics Compute Facility (GPCF))
- Virtualization – Development, Integration, Testing necessary to support the evolving needs of virtualization in a grid/cloud environment. (FermiGrid, FermiCloud, GPCF)
- Cross Department Interest and Involvement in all of these
- More details in FermiGrid presentation
- Plans are being developed to implement and deploy stageable, scalable facilities
  - Subject to priorities, budget, other resources

# Current Login, Analysis and Test facilities for non-Run2, non-CMS experiments (& some Projects)




# Proposed Facilities for FY10



# What happens if we don't build FermiCloud?

- Need place for developing/integrating/testing production level software.
  - Some planned uses of FermiCloud not being done now for lack of resources.
  - Some using the shrinking Gridworks Test facility now.
    - Needs to be moved from FCC.
    - Moving some of the old nodes is not possible, they do not survive power cycling.
  - Some uses are new (e.g. gpcf, storage investigation, virtualization investigation, cloud investigation), no resources currently exist.
- When developing for a highly distributed environment, integration and testing on a laptop does not work.
- ITIL/ISO20000 requires more testing.



*CD FY10 Budget and Tactical Plan Review*

**FY10 Tactical Plans for**  
***Grid/FermiGrid***

FY10 Tactical Plan for FermiGrid	CD DocDB #3330
----------------------------------	----------------

Keith Chadwick

02-Oct-2009

# FY10 Tactical Plan for *Grid/FermiGrid*

- Tactical Plan Leader: *Keith Chadwick*

## Service Activity List

- *Grid/FermiGrid/Operations*
- *Grid/FermiGrid/OSG Services*
- *Grid/FermiGrid/User Support*
- *Grid/FermiGrid/Security*

## Project Activity List

- *Grid/FermiGrid/Development*
- *Grid/FermiGrid/Integration*
- *Grid/FermiGrid/Management*
- *Grid/FermiGrid/Cloud Computing*
- *Grid/FermiGrid/Storage*
- *Grid/FermiGrid/Virtualization*

Note – Red = New or updated activity for FY2010

# Service Activity: *Grid/FermiGrid/Operations*

- Goals Related to this Activity:
  - Ongoing operations of the FermiGrid Services including High Availability.
  - Maintenance of the FermiGrid Services infrastructure.
  - Response to incidents and service requests to assure compliance with OLA's and SLA's.
- Key Metrics:
  - Successful use of the FermiGrid infrastructure by Fermilab personnel, OSG collaborators and LHC experiments.
  - Number of (automatic and manual) interventions required to maintain the infrastructure availability at Service Levels specified in formal OLA's and SLA's.
  - Collection and publication of individual service metrics:
    - <http://fermigrid.fnal.gov/fermigrid-metrics.html>
- Milestones:
  - Deployment of Kernel security updates.
  - Deployments of new Grid Middleware (VOMS, GUMS, SAZ, etc.)
  - Deployments of new versions of Condor and VDT.
- Service Documentation:
  - <http://fermigrid.fnal.gov>
- Issues and Risks:
  1. Continued high level of security patching may take away from other tasks.
  2. We have an ongoing risk to the infrastructure caused by users exploring new areas of the “computing phase space” that result in various overloads or denials of service.



# Service Activity: *Grid/FermiGrid/OSG Services*

- Goals Related to this Activity:
  - Ongoing operation of the OSG Gratia infrastructure.
  - Ongoing operation of the OSG Resource Selection Service (ReSS).
  - Ongoing operation of the OSG VOMRS/VOMS.
  - Ongoing operation of the OSG to TeraGrid Gateway (both GT2 and GT4).
- Key Metrics:
  - Successful use of the OSG Services by OSG management and collaborators.
  - Number of (automatic and manual) interventions required to maintain the infrastructure availability at Service Levels specified in formal OLA's and SLA's.
- Milestones:
  - Production deployment of new hardware to host the OSG Gratia repositories.
  - Timely deployment of new OSG Service software releases (following the FermiGrid Software Acceptance Process as documented in CD DocDB 2684-v4).
- Service Documentation:
  - <http://fermigrid.fnal.gov>
  - <https://twiki.grid.iu.edu/twiki/bin/view/Main/WebHome>
  - <http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=3356>
- Issues and Risks:
  1. Failure to meet the OSG service delivery requirements will likely impact the reputation of Fermilab/FermiGrid.

# Service Activity: *Grid/FermiGrid/User Support*

- Goals Related to this Activity:
  - Ongoing support for the existing FermiGrid user community and new community outreach.
  - FermiGrid Users Meeting
  - Assist new users and groups...
- Key Metrics:
  - User satisfaction
  - Number of service desk tickets submitted.
  - Number of service desk tickets resolved.
- Milestones:
  - New FermiGrid School
  - Development of new classes for FermiGrid School
- Service Documentation:
  - Doc?
- Issues and Risks:
  1. We have a ongoing risk to the infrastructure caused by users exploring new areas of the “computing phase space” that result in various overloads or denials of service.

# Service Activity: *Grid/FermiGrid/Security*

- Goals Related to this Activity:
  - Ongoing security activities associated with the operation of FermiGrid.
  - Plan, coordinate and participate in the Fermilab Open Science Enclave (OSE) Working Group.
  - Maintain the associated infrastructure and evolve it as necessary.
- Key Metrics:
  - Number of Grid security Incidents (fire drill and real).
  - Response time to address Grid security incidents
- Milestones:
  - Updates to the Open Science Baseline.
  - Release of the Grid Trust document and acceptance by CSexec.
- Service Documentation:
  - <https://twiki.grid.iu.edu/bin/view/Security/WebHome>
  - <http://cdorg.fnal.gov/grid/ose/index.html>
  - <http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=2573>
- Issues and Risks:
  1. Failure to adequately support the Grid security activities may result in potential security incidents not being addressed adequately or in a timely fashion, leading to unnecessary compromise of additional systems.
  2. Failure to adequately support the Grid security activities may result in additional DOE oversight.
  3. Failure to adequately support the Grid security activities may result in other OSG or EGEE sites being unwilling to exchange Grid workloads with Fermilab and Fermilab hosted Virtual Organizations.

# Project Activity: *Grid/FermiGrid/Development*

- Goals Related to this Activity:
  - Ongoing maintenance and development of the FermiGrid infrastructure and Grid middleware.
- Key Milestones:
  - Development of Gatekeeper-HA (integration of DRBD with Xen);
  - Development of MyProxy-HA (integration of DRBD with Xen);
  - Development of automated /etc/passwd, /etc/group, /etc/auto.\* files, .k5login files;
  - Deployment and evaluation of “NFS-Lite” gatekeeper on FermiGrid **test** infrastructure;
  - Investigation of automated provisioning (rocks or cfengine) together with automated post-install configuration;
  - Installation of Sun Grid Engine (SGE) batch system on test infrastructure;
- Metrics:
  - Deployment of the above technologies in the production infrastructure.
- Project Documentation:
  - <http://fermigrid.fnal.gov>
  - <https://twiki.grid.iu.edu/twiki/bin/view/Main/WebHome>
- Issues and Risks:
  1. The integration of DRBD with Xen has been more difficult than was originally anticipated and we have had less effort to investigate it over the past year due to other higher priority incidents (Kernel upgrades and BlueArc performance).
  2. Failure to investigate, develop and deploy these technologies will likely result in additional downtime.
  3. DRBD may not work as expected, and we may be forced to investigate alternative clustered filesystems.

# Project Activity: *Grid/FermiGrid/Integration*

- Goals Related to this Activity:
  - Participation in the OSG ITB;
  - Timely integration of newly released Grid middleware into the FermiGrid infrastructure.
- Key Milestones:
  - Completion of Grid middleware and other software integration cycles (following the FermiGrid Software Acceptance Process as documented in CD DocDB 2684-v4).
- Metrics:
  - Number of interventions (incidents and/or service requests) that are required to deploy newly released software into production operations.
- Project Documentation:
  - <http://fermigrid.fnal.gov>
  - <https://twiki.grid.iu.edu/twiki/bin/view/Main/WebHome>
- Issues and Risks:
  1. Lack of effort in the software integration phase makes it more likely that issues will not be caught in integration and they will only be found once the production OSG release has occurred.

# Project Activity: *Grid/FermiGrid/Management*

- Goals Related to this Activity:
  - Ongoing management of the FermiGrid Services project.
- Key Milestones:
  - Timely deployment of new software releases (following the FermiGrid Software Acceptance Process as documented in CD DocDB 2684-v4).
  - Negotiation of OLA's and SLA's with service providers and service clients.
- Metrics:
  - Effective and responsive management of FermiGrid. Measured by the number of times that someone is in line management offices complaining about the operations of FermiGrid.
  - Number of OLA's and SLA's.
  - Performance of FermiGrid against the OLA's and SLA's.
- Project Documentation:
  - <http://fermigrid.fnal.gov>
- Issues and Risks:
  1. Effective management of FermiGrid is a requirement for successful operation of FermiGrid.

# Project Activity: *Grid/FermiGrid/Cloud Computing*

- Goals Related to this Activity:
  - Production deployment of the FermiCloud Infrastructure.
- Key Milestones:
  - Commissioning of the FermiCloud development infrastructure.
  - Commissioning of the FermiCloud production infrastructure.
- Metrics:
  - Number of (configured, running) virtual machines in the cloud infrastructure.
  - Utilization of the cloud infrastructure.
- Project Documentation:
  - <https://cd-docdb.fnal.gov:440/cgi-bin/ShowDocument?docid=3302>
- Issues and Risks:
  1. Developers in the Grid Facilities department (as well as a significant and growing number of Computing Division departments and activities) need access to a rapidly deployable heterogeneous infrastructure (different departments have different needs).
  2. The current FAPL systems in the GridWorks cluster are on their last leg (!) and are largely no longer adequate hosts for software development;
  3. The design of the General Purpose Computing Facility (GPCF) makes extensive use of virtualization and is expected to rapidly evolve in the direction of “Cloud Computing”, the development and integration testing for the GPCF would be done on the cloud infrastructure;
  4. Investigation of alternate storage services (such as Lustre and/or Hadoop) for the Grid would benefit from the Cloud Computing initiative;
  5. To address the above issues, as well as position the CD to take advantage of future developments in Cloud Computing (aka “Infrastructure as a Service”), a dedicated on-site Cloud infrastructure is required.

# Project Activity: *Grid/FermiGrid/Storage*

- Goals Related to this Activity:
  - Ongoing enhancement projects to the FermiGrid production storage deployment.
  - Participation in the investigations of potential Lustre and Hadoop filesystem deployments within FermiGrid to address the high bandwidth needs of the user community.
- Key Milestones:
  - Development, integration and commissioning of Grid storage deployments at the scales required by the FermiGrid user community and agreed to by the Computing Division.
  - Evaluation (and if appropriate deployment) of NFS-Lite for FermiGrid gatekeepers.
- Metrics:
  - Total number of storage services/systems deployed and their size (GigaBytes/TeraBytes).
  - Measured I/O rates to deployed storage services.
  - Measured storage service availability.
  - Reduce or eliminate the impact of user usage of these filesystems on other FermiGrid services.
- Project Documentation:
  - This is a new activity category to capture the effort spent working on storage related issues.
- Issues and Risks:
  1. A small segment (at the moment) of the user community have needs to run I/O intensive Grid jobs.
  2. The current Grid infrastructure is optimized for computationally intensive Grid jobs.
  3. There is a user education issue.
  4. Failure to make the necessary investment in the FermiGrid storage deployments is likely to result in incidents when the user community attempts to run I/O intensive Grid jobs on the current infrastructure that is optimized for computationally intensive Grid jobs.



# Project Activity: *Grid/FermiGrid/Virtualization*

- Goals Related to this Activity:
  - Develop and deploy enhancements to the production deployment of the FermiGrid virtualization infrastructure.
- Key Milestones:
  - Development of tools and/or interfaces for management of the fleet of virtual machines within FermiGrid.
- Metrics:
  - Number of (configured, running) virtual machines.
  - Number of the virtual machines deployed and managed by “higher level” tools.
- Project Documentation:
  - <http://fermigrid.fnal.gov/fermigrid-systems-services.html>
  - <https://cd-docdb.fnal.gov:440/cgi-bin/ShowDocument?docid=3292>
  - <https://cd-docdb.fnal.gov:440/cgi-bin/ShowDocument?docid=3302>
- Issues and Risks:
  1. Failure to invest in these tools/enhancements will result in higher personnel costs to maintain the FermiGrid infrastructure.
  2. Failure to invest in these tools/enhancements will negatively impact the FermiGrid Cloud Computing effort.

# Ripple Effect on Shared IT Services

(What new requirements does your service have for other services)

Activity Level 2	Network Connectivity: Expanded Service	<--	Network-Attached Storage: Add BlueArc Space	<--	GCC Space	<--	Central Services	<--
Storage			1 "rack" of BlueArc	ATS				
Networking	Enhance CDF, D0 & GP Grid Network at GCC	ATS						
Cloud Computing (FermiCloud)	6509 Blade for Cloud Workers	NMP			2 racks	NMP		
OSG Ticket Exchange							FootPrints to Remedy Bridge	ATS
Authentication							Kerberos Authentication	ATS

- FermiGrid has a request to acquire higher performance “dedicated” BlueArc storage (dual raid controllers + high performance disk) to retire ~4 year old equipment and assure “guaranteed” I/O bandwidth (not shared with raid controllers used by other BlueArc clients).
- FermiGrid requests that site networking undertake the necessary infrastructure investments and changes at GCC to support improved network bandwidth and connectivity across the CDF, D0 and GP Grid clusters.
- FermiGrid will deploy (if funded) two racks of Cloud Computing at GCC in FY10.
- FermiGrid relies on the OSG Footprints to Fermilab Service Desk Remedy ticket exchange system.

# FY10 FTE: Request for *Grid/FermiGrid*

Staff	Role	Effort Percentage	
Berman, Eileen	Management	30%	
Chadwick, Keith	Management	40%	
	FermiGrid Services	20%	
	OSG Services	5%	
Green, Chris	OSG Services	20%	OSG Funded
Lowe, Faarooq (just hired)	FermiGrid Services	90%	
	OSG Services	10%	
Sharma, Neha	FermiGrid Services	25%	
	OSG Services	5%	
Timm, Steve	Management	10%	
	FermiGrid Services	80%	
	OSG Services	5%	
Yocum, Dan	FermiGrid	75%	
	OSG Services	20%	
	GridWorks Coordinator	5%	
*** New Hire	FermiGrid Services	90%	
	OSG Services	10%	

# FY10 FTE: Discussion for *Grid/FermiGrid*

- Progress on FermiGrid remains effort limited;
  - Faarooq Lowe has just started (replaces FTE “lost due to retirement” in January 2008).
- FermiGrid expects to host a system administrator “intern” later this (fiscal) year;
  - This will initially be a net loss of effort.
- There are several initiatives that would benefit **(would allow these activities to proceed with greater priority)** from additional effort:
  - General Purpose Computing Facility (GPCF);
  - FermiCloud;
  - Virtualization within FermiGrid and FermiCloud:
    - Pre installation automated configuration;
    - Automated provisioning;
    - Post installation automated configuration;
  - Additional batch systems (PBS-HA, SGE, SGE-HA);
  - Grid storage system investigations and advanced deployments;
  - Automate, automate, automate!

# FY10 M&S: Request for *Grid/FermiGrid*

Qty	Description	Unit Cost	Extended Cost	Fund Type
12	Domestic Travel	\$2,500	\$30,000	Op
6	Foreign Travel	\$3,500	\$21,000	Op
4	New desktops	\$3,000	\$12,000	Op
1	Miscellaneous Operating Expenses	\$10,000	\$10,000	Op
1	Maintenance	\$5,000	\$5,000	Op
1	Training & Documentation	\$10,000	\$10,000	Op
2	Racks of Systems for FermiCloud	\$125,000	\$250,000	Eq
1	6509 blade for FermiCloud	\$10,000	\$10,000	Op
1	BlueArc storage (7 shelves + 2 raid controllers)	\$140,000	\$140,000	Eq
15	Memory & Disk upgrades for CDF, D0, GP Grid Gatekeepers	\$1,500	\$22,500	Op
1	Local high performance storage appliance for FermiGrid to support Gatekeeper-HA and MyProxy-HA	\$50,000	\$50,000	Op
2	Additional core FermiGrid systems	\$8,000	\$16,000	Op
3	Head Nodes / Servers for GP Grid cluster 2	\$5,000	\$15,000	Op
1	Hiring Fee for Posted Opening 090060	\$30,000	\$30,000	Op
1	Hiring Fee for New Hire	\$30,000	\$30,000	Op

# FY10 M&S: Discussion 1 for *Grid/FermiGrid*

- Travel:
  - Request same as last year;
  - So far the indications are that we may be traveling a bit more than last year;
- The FermiGrid BlueArc storage request is in the Central NAS worksheet;
- The need for memory and disk upgrades on the Gatekeepers will be driven by the decision to deploy “NFS-Lite”:
  - In order to address BlueArc I/O issues associated with Grid jobs home area usage;
- We may be able to skip the “local high performance storage appliance”;
  - **IF** we are able to successfully demonstrate the integration of DRBD with Xen;
- Additional core FermiGrid systems:
  - **May** be required to support additional “core” Grid services;
- Head Nodes/Servers for GP Grid Cluster 2:
  - Will be required when the FY2010 worker nodes are purchased;
- Hiring fees:
  - Fee for 090060 **may** have been processed at the very end of FY2009;
  - Need for the other Hiring Fee is dependent on allocation of a posted opening.

## FY10 M&S: Discussion 2 for *Grid/FermiGrid*

- **FermiCloud** is needed to support Grid department developers;
- **and** needed to replace aging FAPL systems in the GridWorks cluster that are failing;
- **and** needed to support Grid storage investigations;
- **and** needed by other departments and groups in the Computing Division and external collaborations (Open Science Grid);
  - Grid/FGS, Grid/OSG, DMS, FEF, HPPC, CET, CMS T1, REX, SSE, GPCF;
  - OSG Storage, OSG + LHC Tier 3, CMS, ATLAS.
- FermiCloud will allow better efficiency and agility:
  - Rather than dedicating physical machines for development that would use power and cooling 24x7, the virtual machines within the cloud can be “shelved” and other virtual machines run on the hardware when the developers don’t need their systems running (i.e. nights and weekends);
  - Rather than waiting for a full hardware acquisition cycle, when new needs are identified new virtual machines can be rapidly deployed once the need is identified, and migrated to the new hardware once it has been delivered.
- The FermiCloud initiative is requested to be funded @ \$250K (two racks), but it can be scaled back to one rack (or even less) based on the available funding.

# Impact of Preliminary Allocation

- ???




# Summary of Past Action Items

- CD Action Item 69 – “Can we develop a plan to host interns for site admin training?”
  - FermiGrid expects to host a site administrator intern later in FY2010.
- CD Action Item 100 - “Is there any connection between this "cloud computing" and other groups?”
  - FermiGrid has met and continues to meet with potential FermiCloud stakeholders.
  - A document detailing the design and current state is available in DocDB:
  - <https://cd-docdb.fnal.gov:440/cgi-bin/ShowDocument?docid=3302>
- CD Action Item 207 – “BlueArc/SAN Discussion”
  - FermiGrid continues to meet and coordinate with Central NAS administrators.

# Tactical Plan Summary – *Grid/FermiGrid*

- FermiGrid made significant advancements while facing significant challenges during FY2009:
  - FermiGrid-HA has demonstrated significant ROI;
  - FermiGrid GUMS, SAZ and Squid services were transitioned to 24x7.
  - BlueArc I/O issues have been identified;
  - See our tactical plan for the complete list.
- We have a new employee:
  - But I'd still like another;
  - See the list of tasks that would benefit from this back on page 20.
- We have significant challenges for FY2010:
  - The FermiGrid roadmap is not as clear as last year;
  - We are actively working to perform the necessary foundational investigations to identify the best approaches to the issues that we face;
  - FermiCloud is a **necessary** investment in the future.



*CD FY10 Budget and Tactical Plan Review*

**FY10 Tactical Plans for**  
**Grid Services**

FY10 Tactical Plan for Grid Services	DocDB#3335

Gabriele Garzoglio

Oct 2, 2009

# FY10 Tactical Plan for Grid/Grid Services

- Tactical Plan Leader: [Gabriele Garzoglio](#)

## Service Activity List

### ***Maintenance and Support for lvl2 / lvl3 activities:***

- Authorization
- Authorization/SAZ \*
- WMS/ReSS
- WMS/GlideinWMS \*
- Accounting
- Security

## Project Activity List

### ***Development and Management for lvl2 / lvl3 activities:***

- Authorization/SAZ \*
- WMS/ReSS
- WMS/GlideinWMS \*
- WMS/GlideinWMS/Corral
- Information Systems
- Accounting
- Metrics Management/MCAS \*
- CEDPS/Storage
- CEDPS/MCAS

# Project and Service Activities under WMS/ GlideinWMS/ (I)

- **Goals Related to this Activity**
  - Provide **maintenance** and **support** for the Glidein Workload Management System (**GlideinWMS**) for CMS, FermiGrid, CDF, OSG, and other stakeholders. Enhance and **further develop** **GlideinWMS** based on stakeholder input.
  - In the context of GlideinWMS, package **gLexec** authorization software in collaboration with gLite developers along with OSG-specific components. Provide **maintenance** and **support** to OSG.
  - As part of the **CorralWMS** project, develop **GlideinWMS** so that it is fully compatible and **integrated** with the **Corral** frontend. Design and develop an application-level monitoring system. Establish collaborative workspaces, training methods, and contribute to project leadership.
- **Key Milestones (Projects)**
  - User-configurable aggregation of Corral/GlideinWMS monitoring and logging data (2nd Q10)
  - Initial release of integrated Corral/GlideinWMS system accepted by existing user communities (4<sup>th</sup> Q10).
  - Future GlideinWMS project releases encompassing new requirements as defined by stakeholders
- **Key Metrics (Services)**
  - Number of trouble tickets reported and resolved for the system.
  - Turn-around time in resolving incidents.
  - Number of jobs executing in parallel on CorralWMS and GlideinWMS (goal: >20,000).
  - Number of jobs per day on CorralWMS and GlideinWMS (goal: >100,000).
  - Number of user communities for CorralWMS and GlideinWMS (goal: >4).

# Project and Service Activities under WMS/ GlideinWMS/ (II)

- Project Documentation:  
<http://www.uscms.org/SoftwareComputing/Grid/WMS/glideinWMS/>
- Issues and Risks (specific to this activity, includes allocation impact)
  - Failure to operate and support GlideinWMS may have a major effect on the efficiency of ongoing data analyses for CMS, CDF, and the Minos experiment. Failure to respond to new stakeholder requests may delay deployment by FermiGrid and OSG, effectively lowering efficiency across those cyberinfrastructures.
  - Failure to package new gLexec releases for OSG may prohibit wide-spread use of pilot-based workload management systems (GlideinWMS, ATLAS's PANDA) on OSG, lowering efficiency and possibly decreasing the pool of potentially available resources.
  - Failure to meet the CorralWMS objectives would have a negative effect on the spread of pilot-based workload management systems and increase the cost for organizations wishing to use both OSG and TeraGrid. As the project is externally funded, failure to make progress in FY10 could also impact the funding of the effort.
  - Development infrastructure is getting old and risks forced decommissioning: the project wants to use 3 VM from the FermiCloud for testing.

# Project Activities under

## Metrics Management/MCAS/ (I)

- **Goals Related to this Activity [from Tactical Plan]**
  - **Develop** and **deploy** a Metrics Analysis and Correlation Service (**MCAS**). Focus on US CMS, RunII, and the future neutrino experiments.
  - The MCAS Project provides an integral solution for system operators or VO users to uniformly access, transform, and represent disjoint metrics data generated by distributed middleware or user services. The infrastructure assists indexing, navigation of existing metrics as well as it supplies tools and services to define and store quantifiable on-demand document data. The project basic interfaces will be developed to satisfy criteria of low cost integrability into existing applications. Simply put, the project offers users a simple web based toolkit for organizing and plotting any kind of system state data.
- **Key Milestones**
  - Minos portal (Sep 09);
  - Data source administration tool. Data cache to production (Nov 09);
  - Data transformation and business analysis layer. Data cache operations support tools (Dec 09);
  - User friendly data analysis/query front end (Lower priority; Mar 10);
  - Data schema browser (Lower priority; Apr10);
  - Additional operational tools. Interface with Enterprise-level reporting tools (May 10).
  - New rendering primitives. Feature adjustment/support (ongoing).
  - Monthly meetings with each CDF, DZero, CMS, Minos stakeholders individually.
  - Bi-monthly meetings with other groups monitoring the project to resolve overlap in functions.
  - Milestones status: <http://sourceforge.net/apps/trac/mcas/report/1>

# Project Activities under Metrics Management/MCAS/ (II)

- Project Documentation:  
<http://www.fnal.gov/docs/products/mcas>
- Issues and Risks (specific to this activity, includes allocation impact)
  - The two key features of the MCAS project are
    - providing quick and easy access to building custom user-level dynamic reports
    - documenting the schema of user/system metrics

Lack of common services to generate reports on metrics will encourage users to do uncoordinated development/integration to achieve objectives similar to the ones of MCAS. These activities will inevitably lead to duplication of effort and potential software vulnerabilities, due to the lack of common development standards. Failure to formalize and catalog metrics will hinder the ability to report the performance of computing services operations to stakeholders, including to CD management.

- The project uses desktops and laptops as main development platforms. This limits the ability to test the integrated system: the project asks at least 1 VM for development / testing from the FermiCloud.



# Project and Service Activities under Authorization/SAZ/ (I)

- Goals Related to this Activity [from Tactical Plan]
    - **Develop** and **maintain** the **SAZ** service to enable user/vo/role/ca banning on campus grid facilities, in particular on FermiGrid, and to provide support to customers of the SAZ software.
  - Key Milestones (Projects)
    - User input validation; simplify / modularize code \*\*; improved DB connection management; rewrite calls to shell commands using APIs (Oct 09).
    - Address Resource Exhaustion for Sockets and Threads \*\* (Nov 09).
    - Integrate XACML call-out protocol \*\* (Dec 09).
    - Address OSG user-banning requirements (Spring 10).
    - Quarterly internal meetings
- \*\* Development items to improve response of the SAZ server to user “storms”. Operational items include segmentation of domains (so that a storm affects only the originating domain) and job submission rate limitation (in the next Condor release)

# Project and Service Activities under Authorization/SAZ/ (II)

- Key Metrics (Services)
  - Number of users banned through SAZ.
  - Number of trouble tickets for SAZ.
  - Number of recommended security improvements implemented
  - SAZ on-line metrics:
    - <http://fermigrid.fnal.gov/metrics/fermigrid-ha-saz-data.html>
    - <http://fermigrid.fnal.gov/monitor/fermigrid-ha-saz-ha-monitor.html>
- Project Documentation: <http://saz.fnal.gov>
- Issues and Risks (specific to this activity, includes allocation impact)
  - Failure to maintain and support the SAZ service will result in increased operational complexity in reacting to security incidents for the FermiGrid facility.
  - The development infrastructure for SAZ is old and risks forced decommissioning: the project wants to use at least 1 machine from FermiCloud.

# Goal Description of remaining Tier 2 Activities (I)

- **Authorization**
  - **Maintain** the infrastructure for VO membership registration, focusing on the **convergence of VOMRS and VOMS-admin**. Investigate new mechanisms for **VO and site policy** definition, publication, and enforcement.
- **Accounting**
  - Work closely with stakeholders to identify and appropriately prioritize their **maintenance** needs from the **Gratia** software stack (text and graphical reports, probes and collectors).
  - Provide **expertise** and **code updates** as needed by the groups **operating** the production instance of **Gratia Collectors** within OSG and the Fermilab Computing Division.
  - Ensure that (potential) **Gratia Extensions** provided by external projects are **coordinated** and well integrated into the existing **code base**, test, release and support mechanisms. Improve quality assurance processes for the software.

# Goal Description of remaining Tier 2 Activities (II)

- **WMS/ReSS**
  - **Finish development** activities for the Resource Selection Service (**ReSS**) project. Move all software to maintenance. Close the project. **Provide** second-level **support** to the FermiGrid operations of the service for OSG and CD.
- **Information Systems**
  - Contribute to the **maintenance** and **support** of the OSG grid information system (**GIP**). **Begin development** to provide GLUE 2.0 compliance.
- **CEDPS**
  - [NOTE: The CEDPS project is undergoing a restructuring and current funds are blocked. The following goal for the activity may need to be revised.]  
Integrate GlideinWMS with the DataKoa data management service, interfacing with the Magellan cloud infrastructure.
- **Security**
  - Perform **security-focused reviews** of several software projects.

# Ripple Effect on Shared IT Services

(What new requirements does your service have for other services)

Activity Level 2	FermiGrid	<--
WMS	3 VMs + 2 GridWorks Repl. on FermiCloud	TA
Authorization	>1 VM + 1 GridWorks Repl. on FermiCloud	TA
Accounting	3 VMs + 10 GridWorks Repl. on FermiCloud	TA
Metrics Mngmt	>1 FermiCloud VM	TA
Metrics Mngmt	2 Prod / Int. Servers	MA

- In the Grid Services budget, **we rely on the FermiCloud infrastructure** for WMS, SAZ, Gratia, and MCAS development / testing machines (see FermiGrid budget)
- Number of VM is bare minimum. We envision that GridWorks machines for devel and testing will keep dying and will need replacements.

# FY10 FTE and M&S: Request

Level 0/1 Activity: Grid / Grid Services (CD Funds)

Activity Level 2	Project or Service	Project Priority	FY10 FTE Request At Activity Level 2	FY10 M&S Request At Activity Level 2
Metrics Managmnt	Project	Medium	1.6	15000 (Prod and Int servers) + >1 VM (devel)
WMS	Proj/Srv	High	2.3 (0.4 from CD)	3 VM (devel / test) + 2 VM (GridWorks repl.)
Authorization	Proj/Srv	High	0.5	>1 VM (devel / test) + 1 VM (GridWorks repl.)
Accounting	Proj/Srv	High	0.6 (0.35 from CD)	3 VM (devel / test) + 10 VM (GridWorks repl.)
CEDPS	Project	Medium	0.6 (0 from CD)	-
Information Systems	Project	High	0.3 (0 from CD)	-
Group Travel (International)	Project			5 x 3500 = 17500
Group Travel (Domestic)	Project			8 x 2500 = 20000
CEDPS Travel (Domestic)	Project			1 x 2500 = 2500
Group Training and Documentation				15 x 1000 = 15000
Group Lap Tops				2 x 2800 = 5600
Group Misc. Operating Exp.				1x5000 = 5000
<b>Total</b>			5.9 (2.85 from CD)	\$80,600

- Requests for VMs are NOT budgeted in the Grid Services M&S: see FermiGrid budget.

# Summary of Past Action Items



- None

# Tactical Plan Summary

- Grid Services plans for activities in these domains: Authorization, WMS, Information Systems, Accounting, Metrics Management, Security, and CEDPS.
- These are funded by 5 sources: CD, OSG, CMS, SciDAC2, STCI.
- The personnel comes from the OSG group, FermiGrid, CMS Grid Services, Physics Software Tools.
- We rely on the FermiCloud initiative to fund development and test machines for the Grid Services projects. The GridWorks cluster for development and testing will need replacement soon.





*CD FY10 Budget and Tactical Plan Review*  
**FY10 Tactical Plan for**

Open Science Grid at Fermilab	CD-3326
-------------------------------	---------

[Ruth Pordes, covering input from Mine Altunay, Chris Green,  
Tanya Levshina, Dave Ritchie, Chander Sehgal, ]

2<sup>nd</sup> Oct 2009

# FY10 Tactical Plan for GRID/Open Science Grid



## Service Activity List

- OPEN SCIENCE GRID/COORDINATION
- OPEN SCIENCE GRID/COMMUNICATIONS
- OPEN SCIENCE GRID/SECURITY
- OPEN SCIENCE GRID/USER SUPPORT
- OPEN SCIENCE GRID/VDT

## Project Activity List

- OPEN SCIENCE GRID/SECURITY\_RESEARCH

# FY10 FTE and M&S: Request

Level 0/1 Activity: [GRID / Open Science Grid](#)


Activity	Priority	Total M&S	Total FTE		Base M&S	Base FTE		External M&S	External FTE
Communication	High	12,500	0.35		7,500			5,000	0.35
Coordination	High	26,000	1.64		6,000			20,000	1.64
Security	High		1.06			0.21			0.85
User Support	High		0.67						0.67
VDT	High		2.06			0.06			2
		85,000	0			0		85,000	0
<b>Total</b>		<b>123,500</b>	<b>5.78</b>		<b>13,500</b>	<b>0.27#</b>		<b>110,000</b>	<b>5.51</b>
* Reflect \$70K reduction in M&S already communicated to Financial Team									

Additional 0.3 FTE +20K M&S provided for iSTGW editor from external funds  
# in administrative and web support.

# Current OSG Project Budget Profile

## Rampdown during 2010-2011

### OSG Budget Profile



# External Project funds to Fermilab

OSG Funds to Fermilab	NSF- arrives at Start of Fiscal year	DOE allocated 4.5 Years of Funding which Arrives ½ way through prior fiscal year	Total
FY10	\$590K	\$873K	\$1,463K
FY11 – allocated to date	\$590K	\$449K	\$1,039K

- Fermilab will need \$400K reallocated from other funds to receive full funding for FY11. We believe the \$ exist in the project as a whole but there are no plans yet in place to do this.
- Discussions of whether/how to distribute remaining funds for FY12 if continued funding not in place yet to happen.

# OSG Future Planning as agreed by the OSG Council in Mar and Aug

## Phase 1 by June 2009

- Gather and document requirements & expectations from major stakeholders for 2010 to 2015
- Solicit guidance from OSG Council on key directions for future

## Phase 2 by Aug 2009

Done

- Face-to-Face Council Mtg
- 2-page plan for each OSG work area (or functional unit) **Some 2 page strategy documents written. Summary of status to follow.**
- List/abstract of satellite proposals **More Satellites being proposed and some funded. More to Follow**

## Phase 3 by Dec 2009 - in progress

- Analysis => outline for proposal
- Document Architecture
- Identify participating senior personnel (and institutions)

## Phase 4 at March 2010 All Hands Meeting

- Endorsement of proposal by stakeholders

# Thoughts on new proposals

- Guidance to submit separate proposals to DOE and NSF. SciDAC II cycle up at the end of FY11. experience last time round was some months gap between SciDAC I and II funding.
- Proactive work by NSF to get OSG and TeraGrid to align – meeting of Jot and project leads now expected at DOE on Nov 5<sup>th</sup>.
- Hope for 5 year proposals augmented by Satellite Projects.
- US LHC remain key accounts for OSG. Their drivers will significantly affect the new proposals.
- OSG remains a strategic direction for US LHC and Fermilab.

# Thoughts on new proposals

- US LHC Input to date:
  - Continue support for Operations, Software, Security
  - Continue interoperation and collaboration with Europeans
  - Address issues of >100 Tier-3s and 15 Tier-2s and promote cross-campus, national and international distributed facility scope.
  - Take ownership of the architecture and respond to scalability, ease of use, data management and access issues.
  - Be prepared to help off-shore sites on request – to date especially in Brazil.
  - Participate in Satellite Projects for evolution of s/w etc.

-



# Thoughts on new proposals

- Fermilab strategy for continued contributions to and external funding from OSG?
  - Bring Security leadership and expertise to the table.
  - Continue as significant part of the project team for Software support and provision.
  - Continued and perhaps expanded support for Operations especially for US LHC.
  - Continued leverage of software development and evolution/collaboration for Information Services (GIP, RESS), Workload Management (GlideinWMS), Accounting (Gratia), Production Facility Services (ITIL, SLAs, Clouds, testing environments etc). Participate in Satellite Projects for collaborative funding of evolution of s/w.
  - Increase leverage across OSG @Fermilab and US CMS (& US ATLAS) support (all Tiers, S&C operations, evolution of worldwide distributed infrastructure)
  - Continued and perhaps expanded User Support for all/new Fermilab user communities, other DOE users (e.g. ITER) and use of other DOE facilities (e.g. LQCD use of ACLF)
  - Project Staff Office linked to Executive Director role. Council sub-committee to address length of term. Expect recommendation that current term last til new funding in place + 1 year with Deputy assigned for final year and then transition.

# Service Activity: OPEN SCIENCE GRID/COORDINATION

- Goals Related to this Activity [from Tactical Plan]
  - Meet deliverables of the OSG WBS as defined in the Fermilab OSG SOW
    - Work of the OSG Executive Director to deliver her responsibilities to the OSG Project.
    - Work of the Project Manager to deliver his responsibilities to the OSG Project
    - CD administrative support for OSG
    - Work of CD management in coordinating CD staff contributions to OSG.
    - NeesComm participation (project advisory committee & community collaboration)
- Key Metrics
  - Success in OSG reviews and attempts to obtain sustainable funding.
- Service Documentation:
  - <https://twiki.grid.iu.edu/bin/view/Management/WebHome>
  - Fermilab SOW FY09 (FY10 in preparation)  
<https://cd-docdb.fnal.gov:440/cgi-bin/RetrieveFile?docid=3027&version=1&filename=year3%20OSG%20FNAL%20SOW%20fully%20signed.pdf>
- Issues and Risks (specific to this activity, includes allocation impact)
  - Failure of OSG to provide production quality services that Fermilab's facilities and the OSG can use.
  - Failure of other institutions within the OSG to deliver their commitments will affect the ability of the project to succeed.
  - Differences in Fermilab and other OSG stakeholder needs and schedules can lead to divergence in implementations and services.

# Service Activity: OPEN SCIENCE GRID/COMMUNICATION

- Goals Related to this Activity [from Tactical Plan]
  - Work of the OSG Communications and Education Coordinator and his staff to deliver their responsibilities to the OSG Project.
    - Support the OSG Web portal content.
    - Attend OSG VO meetings, write minutes and pay attention to ideas for communication that come out of these meetings.
    - Provide updated posters, postcards, brochures and other print media. Publish multi-media content on request
    - Communicate with External Communicators at OSG staff and Consortium institutions via one-on-one email conversations, Quarterly Phone Meetings, Annual Face-To-Face at AHM
    - Produce Research Highlights
    - Produce iSGTW
    - Produce OSG News
    - Coordinate Educational activities within the OSG staff e.g. Joint OSG TeraGrid Student programs at OSG AHM, Joint OSG TeraGrid Student programs at TG10, SuperComputing 09 Education Programs, Grace Hopper
    - Publicise and ensure progress in the OSG Computer Science Fellowship program
    - Publicise OSG intern and summer student activities
    - Contributions of CD staff as part of overhead.
- Service Documentation:
  - <http://www.opensciencegrid.org>
- Issues and Risks (specific to this activity, includes allocation impact)

# Service Activity: OPEN SCIENCE GRID/SECURITY

- Goals Related to this Activity [from Tactical Plan]
  - Manage OSG security team, coordinate activities within OSG, prepare roadmaps and planning
  - Support the needs of T1-T2-T3 sites, provide direct communication to diagnose and solve problems
  - Identify and develop pro-active monitoring tools to prevent incidents
    - Develop a gratia based prototype for early investigation
  - Conduct risk assessment and auditing of OSG assets
  - Contribute to documentation effort
  - Provide training materials and conduct training session for OSG staff, site and VO security contacts
- Service Documentation:
  - <https://twiki.grid.iu.edu/bin/view/Security/WebHome>
- Issues and Risks (specific to this activity, includes allocation impact)

# Service Activity: OPEN SCIENCE GRID/USER SUPPORT

- Goals Related to this Activity [from Tactical Plan]
  - Contributions to OSG Engagement area and other VO software support to deliver responsibilities to the OSG Project.
    - Gratia operations, support and maintenance following OSG requirements
    - Engagement with ITER, Geant4, Cactus, LQCD, NERSC resource, OSG-TG Gateway applications,
  - Contributions by CD to OSG engagement and VO software support.
- Key Metrics
- Service Documentation:
  - <https://twiki.grid.iu.edu/bin/view/Accounting/WebHome>
  - <https://twiki.grid.iu.edu/bin/view/Engagement/WebHome>
- Issues and Risks (specific to this activity, includes allocation impact)

# Service Activity: OPEN SCIENCE GRID/VDT

- Goals Related to this Activity [from Tactical Plan]:: Contributions to OSG VDT to deliver responsibilities to the OSG Project:
  - Make leadership contributions to the better integration of the VDT team across the “non-storage” and “storage” sub-groups.
  - Package, distribute and support storage, and other modules on request, components of the Virtual Data Toolkit.
  - Provide "Level 2" User Support for LHC Tier-2/Tier-3 deployments of OSG Storage Software.
  - Provide support for LIGO, SCEC and other users of storage – both community owned and opportunistically available – on OSG accessible resources.
  - Provide initial support interface and coordination between the users of VDT storage and the development groups - dCache, BeStMan, Xrootd, Hadoop, GridFTP, Gratia Storage and Transfer probes.
  - Participate in WLCG Technical Group work (led by Maarten Litmaath), report issues to OSG EB.
  - Make significant contributions to the OSG Software Tools Group and OSG software coordination.
  - Collaborate with OSG scalability area (UCSD) on storage testing, performance and scalability improvements
  - Work with WLCG storage groups to help provide what the US LHC stakeholders need.
  - Work with the OSG Tier-3 area to understand and proactively support both US ATLAS and US CMS Tier-3 storage and other VDT software needs.
  - Participate in planning storage and other OSG software needs and implementations with new communities
  - Participate in OSG Site Administrator, Tier-3 and other workshops on request, provide storage tutorial, installation tips etc
  - Manage OSG storage group
- Contributions from CD to VDT testing and packaging, e.g. Prima, GLExec work on Wisconsin Build and Test system.
- Service Documentation:
  - <https://twiki.grid.iu.edu/bin/view/Storage/WebHome> <http://vdt.cs.wisc.edu>

# Project Activity: OPEN SCIENCE GRID/SECURITY\_RESEARCH

Activity type: Research. Starting in FY10 with new staff coming 10/23.

Discssion with Security Group to move activity under their tree.

Thinking started on:

- Modeling and Validation of Emergent Behavior under Threat on Multi-Domain Interconnected Systems (proposal submitted to DOE ASCR with ANL)
  - Social networking models for risk analysis of security incidents (work with Maths at ANL)
  - Access control policy analysis: comparison, testing and validation. (Dr.Vincent Hu, NIST)
  - Pattern analysis of distributed log and accounting files (U Notre-Dame?)
  - Data privacy testing
- 
- Metrics: Success in sustainable funding attempts and resulting scientific publications.

# Ripple Effect on Shared IT Services

– no new ones here.

Activity Level 2	From	What
VDT	GRID / FermiGrid / Cloud Computing	Capable and performant test systems for VDT software testing and validation.
Project Office Support	CENTRAL IT SERVICES / Web & Collaboration Services / DocDB CENTRAL IT SERVICES / Web & Collaboration Services / Indico	DocDB and Indico Service Support
Activity Level 1: Open Science Grid	CENTRAL IT SERVICES / ServiceDesk / Operations	Continued production connection between Service Desk and OSG Operations Center/Footprints
	'GRID / FermiGrid / OSG Services	Operations support for OSG services (over and above those funded by the OSG Project)



# Summary of Project Progress

- Successful project review in January 2009. Project encouraged to continue and plan for the future. Also asked to document a “close down” plan.
- Satellite projects with Fermilab involvement, on the table:
  - GlideinWMS – funded
  - DOE Security research – Modeling and Validation of Emergent Behavior under Threat on Multi-Domain Interconnected Systems; expect any announcements soon
  - Gratia extensions – in review
  - OSG-TG joint activities ExTENCI being written; putting in for US CMS application and HPC dept administrative testing environment of Lustre WAN.
  - Possibility – interest at NSF for co-funding for interoperation and collaboration with European futures (e.g. ROSCOE, data management...)

# Tactical Plan Summary

- OSG @ Fermilab continues to significantly benefit from and leverage CD expertise and development.
- OSG @ Fermilab aims to add value to the existing CD activities and contributions to the US LHC and other Fermilab user communities.

# Input to DOE Accomplishments in Aug 2009




# The Tevatron and Open Science Grid Scientists Join Forces in the Search for the Origin of Mass


Enabled by advanced Distributed Computing, Tevatron scientists shrunk the Higgs mass window from a few hundred GeV to ~ 40 GeV. Thus providing **a 135GeV +/- 15% target region for future discoveries.**

Prevailing theory of mass predicts a “Higgs” particle, and all its properties, except its mass!

Science Goal: Reduce the window of allowed Higgs mass via direct searches, and precision measurements (press release March 2009).



Direct Higgs Search at Tevatron now excluding high mass higgs.



Top and W masses now known to 0.3% and 0.75% respectively.

=> Constrain Higgs via Quantum effects.

# Novel Framework Transforms the Power of Distributed Computing Through Overlay Resource Management

US developed  
framework adopted by  
the *Worldwide Large  
Hadron Collider* (LHC)  
collaborations



**Community Specific**

**High Throughput**



**Science  
Applications**

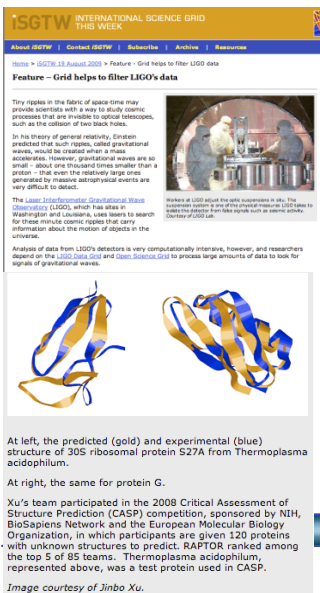
**Efficiency and  
scalability for *high  
throughput* scientific  
applications  
increased by up  
to 10 fold through  
*just in time*  
scheduling.**

**Overlay Resource Manager  
Transforms a Heterogeneous set of  
Distributed Resources into a Uniform  
& Effective Computing Environment**


**Distributed processing and storage resources at  
~60 US Labs and Universities**

# Multi-Science Collaboratory for Data Intensive Science and Distributed Computing

*Collaborative Consortium of domain and computer scientists from more than 100 organizations sustained for more than five years.*



Scientific discovery in a unique national partnership between **computational & computer** scientists who advance the state of the art in distributed computing through leadership in data intensive science.



3D replica of senecavirus, a pathogen discovered several years ago by researchers in Pennsylvania. UC San Francisco researcher Eric Delwart and his colleague Chunlin Wang of Stanford University use the RENCi-developed TeraGrid Science Gateway and the Open Science Grid to access grid computing resources in their search for new viruses.

Image courtesy of the Institute for Animal Health, UK